

MAC–Level Partial Checksum for H.264 Video Transmission over 802.11 Ad Hoc Wireless Networks

E. Masala, M. Bottero, J.C. De Martin*

Dipartimento di Automatica e Informatica / *IEIIT–CNR
Politecnico di Torino, corso Duca degli Abruzzi 24, 10129 Torino, Italy
Phone: +39-011-564-7036 / -5421, fax: +39-011-564-7099
E-mail: [masala | mauro.bottero | demartin]@polito.it

Abstract— We present a technique to enhance real-time video transmissions over 802.11-based ad hoc networks. We propose to limit the coverage of the error detection mechanism to the most error sensitive bit class contained in video packets. Packets are retransmitted only if errors affect those bits, instead of retransmitting all corrupted packets irrespective of error positions. The expectation was that the negative effects of errors in the less sensitive bits should be more than adequately counterbalanced by the lower number of discarded video packets. Moreover, the lower number of retransmissions should reduce network load, with positive effects on all the transmissions. We simulated the transmission of H.264 video over an ad hoc scenario with a varying number of relay nodes using *ns* with experimental bit error traces. Results show consistent video quality improvements (up to 1 dB PSNR) as well as network load reduction with respect to state-of-the-art techniques, such as unequal error protection based on different retry limits.

Index Terms— partial checksum, ad hoc networks, H.264 video transmission, IEEE 802.11

I. INTRODUCTION

Ad hoc wireless networks are increasingly attracting attention because of their ability to connect nodes without relying on pre-existing network infrastructures. At the same time, the growing interest in wireless multimedia applications stimulates the development of techniques to successfully provide those applications over wireless ad hoc networks. Several challenges, however, need to be addressed, such as potentially high error rates and delays.

Many wireless ad hoc networks adopt the widespread IEEE 802.11 wireless local networking (WLAN) standard [1] as the link-layer communication protocol. Several studies focused on performance enhancement of multimedia communications over 802.11 wireless LANs. For instance, layered coding coupled with Unequal Error Protection (UEP) obtained by using different retry limits at the link level has recently been shown to deliver interesting results [2] [3] [4].

Other multimedia transmission enhancement techniques are based on the error resilience features included in most multimedia compression standards. The decoder can often exploit

those features to recover useful information even in partially corrupted packets. The current IEEE 802.11 Media Access Control (MAC) layer, however, prevents the forwarding of erroneous packets, irrespective of error number and position. Corrupted packets are discarded and the sender will retransmit the data until a given maximum retransmission limit (retry limit) is reached [1].

Recent works addressing the issue of partially corrupted packets in wireless multimedia communications propose to use the UDP Lite transport protocol [5]. For instance, in [6], the performance of the UDP Lite protocol is evaluated by means of actual 802.11b experiments using different physical layer transmission speeds. The impact of errors at the application level is considered for the case of video, showing that the quality significantly degrades when the physical transmission speed is increased. In that work, however, the coverage of the UDP Lite checksum is limited to protocol headers, and the 802.11 MAC level error checking feature is completely disabled, hence no MAC level retransmissions are used.

In [7], a modified version of the UDP Lite protocol is proposed. It features a checksum for the packet header, and at the same time it provides an interface to forward all the information supplied by the CRC failures in link-layer frames to the application layer, to improve error location inside the packets. This protocol is combined with a UEP scheme applied to fixed-size link-layer frames in a 3G wireless scenario. Other works suggest to limit the UDP Lite partial checksum to the packet header [8], focusing on schemes that add redundancy at the data link level, and allowing packets containing errors to be forwarded to the applications. However, the quality of the received multimedia streams is not measured.

Recently, the possibility of a partial checksum at the *link* layer has been explored. Link-layer error detection is particularly attractive for moderate or high end-to-end delay scenarios, where end-to-end retransmission schemes are generally not applicable. Hop-based retransmission, in fact, is very fast, delivering “acceptable” packets with lower delay. In [9] [10] [11] a partial checksum approach has been explored for speech and video communication, showing both perceived quality

This work was supported in part by MIUR, Project FIRB PRIMO, <http://primo.ismb.it>

and network performance improvements under several channel conditions.

In this paper we extend our previous work [11] on link-layer partial checksum transmission scheme for video transmission to the case of ad hoc networks, showing the suitability of the approach for this particular type of network. We propose a modification of the IEEE 802.11 link layer to better support video communications in ad hoc networks allowing partially corrupted packets to be forwarded by relay nodes (and not discarded). Video packets are retransmitted by the MAC protocol only if errors are detected in the most sensitive subset of the compressed bitstream, that —if corrupted— would introduce noticeable visual artifacts into the decompressed video sequence. Results obtained simulating an ad hoc network with a variable number of relay nodes show that the negative effects of errors in the less sensitive bits are more than counterbalanced by the lower number of discarded video packets. Moreover, avoiding to retransmit corrupted but “acceptable” packets limits the network load with benefits that are particularly evident in ad hoc networks.

The paper is organized as follows. Section II describes the proposed partial checksum technique for wireless multimedia transmissions and analyzes its advantages. Section III describes an implementation of the partial checksum scheme for the specific case of H.264 video transmission. Section IV explains the experimental setup, including encoder configuration and simulation scenario. Results, including comparisons with reference schemes, are presented in Section V. Conclusions are drawn in Section VI.

II. PARTIAL CHECKSUM FOR WIRELESS MULTIMEDIA

Wireless networks, unlike their wired counterparts, are prone to bit errors. As long as errors affect the less perceptually important data bits, modern multimedia applications can often deal with corrupted frames better than with lost ones by means of error resilient coder design and error concealment techniques. However, if the more perceptually important bits are corrupted, such as header data, the video quality may degrade considerably.

Recent video standards include a number of techniques to enhance the robustness of the compressed data streams [12], for instance the Reversible Variable Length Codes (RVLC) and the insertion of resynchronization markers. Other techniques, such as Data Partitioning (DP) and layered coding, provide an a priori classification and separation of bits according to their sensitivity to errors and losses. These coding modes are usually coupled with Unequal Error Protection (UEP) schemes to take advantage of the different error sensitivity of the various classes or layers.

The UDP Lite protocol [5] is a proposal to introduce a partial checksum at the transport level that protects only the header and the sensitive part of the payload of a UDP data unit. If an error occurs in the checksummed bits, the receiver should drop the packet, otherwise it is forwarded to the application. The usefulness of the transport level partial checksum, however, relies on the possibility to prevent the




Fig. 1. Data-type 802.11 MAC frame structure modified for partial checksum.

MAC layer from dropping corrupted packets. Moreover, partial checksum at the transport level implies that packet checking is performed only at the transmission end points. Therefore, UDP Lite supports only end-to-end error control mechanisms, and it cannot provide fast hop-by-hop retransmission unlike the 802.11 MAC level partial checksum that we propose. The benefit of fast hop-by-hop retransmission is particularly evident in scenarios characterized by high end-to-end delays, in which end-to-end retransmission-based robustness techniques are inapplicable.

In this paper we argue that, for 802.11-based ad hoc networks, a partial checksum approach at the *link layer*, coupled with data partitioning can improve the performance of video transmissions. The fundamental idea is that no error detection needs to be performed on the perceptually less relevant bits, which are forwarded even if corrupted. Application level error resilient techniques are then exploited to recover the maximum amount of useful information from the corrupted data. Note, however, that the MAC level checksum cannot be completely disabled due to the high bit error rate of wireless communications, hence the necessity of a partial checksum mechanism. MAC level headers, in fact, definitely need to be checked for errors to prevent misdelivered packets and to ensure the correct behavior of the 802.11 MAC protocol at each station. Moreover, higher level protocol headers, such as IP, UDP and RTP, need to be protected as well to ensure correct protocol operations.

We propose to design a flexible partial checksum mechanism introducing a two-byte fixed-length field before the payload that indicates how many bits from the beginning of the MAC frame must be considered to compute the frame FCS. The proposed frame structure is shown in Figure 1. In our proposal, UDP Lite is only used to allow a corrupted MAC payload to reach the application level.

Finally, note that the time needed to successfully send a MAC frame can be quite large when compared to the physical layer transmission time of a single MAC frame. A channel contention phase, in fact, is needed for each retransmission, and the contention window doubles after each unsuccessful retransmission. Limiting the MAC layer transmissions as much as possible, as done in the *link layer* partial checksum technique, delivers advantages in terms of increased network throughput and reduced load, which is particularly useful to enhance the performance of ad hoc networks.

III. PARTIAL CHECKSUM FOR H.264 VIDEO TRANSMISSION

In our experiments, the data partitioning functionality of the H.264 standard [13] is employed to pack the bits in sensitivity order. The compressed video stream is subdivided




Fig. 2. Comparison of transmission schemes: proposed partial checksum (a), full checksum (b), header-only partial checksum (c) and UEP (d). Both MAC and IP/UDP/RTP protocol headers are always protected. The shaded area represents the checksum coverage.

by the encoder into three partitions or classes. According to the standard, class A, the most important one, is used for headers, including macroblock headers, and for motion vector information. Class B and C are designed to contain the texture information of the various types of macroblocks. In the proposed partial checksum scheme, all the data belonging to a slice is arranged into a single packet. Class A bits are placed at the beginning of the packet, and their number determines the checksum coverage. The remaining part of the packet is filled with class B and C bits. The proposed packet structure is shown in Figure 2(a). Note that the encoder/packetizer is responsible to set the appropriate value in the two-byte field of the MAC packet that signals the amount of checksummed bits to the MAC layer. For non-video packets, this field is set to the MAC packet length.

To evaluate the performance of the proposed *partial checksum* scheme, Section V presents comparisons with three other transmission schemes. Two schemes, namely the *full checksum* and the *header-only partial checksum*, feature a different amount of checksummed bits per packet. In the full checksum scheme, shown in Figure 2(b), all the bits of each packet are included in the checksum computation, hence that scheme represents the standard 802.11 behavior. The header-only partial checksum scheme, illustrated in Figure 2(c), is based on the same principle of the proposed partial checksum, but it restricts the checksum computation to the bits of the MAC and transport level headers. Finally, an *unequal error protection* scheme, shown in Figure 2(d), has been implemented for comparison purposes. It uses the 802.11 standard checksum mechanism, but the H.264 partitions are placed into two separate MAC frames with different retry limits. Examples of such approach are described, for instance, in [3] [4]. The same retry limit value as in our proposed scheme is assigned to packets containing the H.264 A-type partition, while packets containing the H.264 B- and C-type partitions are sent only once.

Note that only the full checksum scheme is fully compliant with the current 802.11 standard; all the other schemes require either the partial checksum feature or the option to set a different retry limit for each packet class. Both features are not supported by the current 802.11 standard [1]. Proposals such as the 802.11e standard are expected to address some of

these issues.

IV. SIMULATION SETUP

Video sequences at QCIF resolution, 15 fps, are encoded with a fixed quantization parameter using the standard JM encoder, version 6.0a, modified to support standard-compliant data partitioning. The resulting bitrate is about 100 kbit/s. Each frame is subdivided into three slices, each one corresponding to three consecutive rows of macroblocks, thus 45 packets per second need to be transmitted for each flow. The mapping of the partitions into packets depends on the transmission schemes, as shown in Figure 2.

The following concealment techniques have been used. Packet losses are detected by means of the RTP sequence number. In all the schemes under test, we assume that any video unit (i.e. slice) whose class A bits are corrupted or lost cannot be decoded without introducing a large distortion, thus the whole video unit, including class B and C information, is considered lost. In this case, the decoder applies a temporal concealment technique that replaces the missing pixels with the ones in the same position in the previous frame. Header-only checksum technique excepted, class A bits —if received— are always correct because they are checksummed. For the header-only checksum scheme, instead, class A bits could be corrupted. In our implementation class A information is discarded if any syntax violation is detected during the class A decoding process. If the class A bits are correctly received and class B or C bits are available, the JM 6.0a decoder should take advantage of the bitstream error resilience to improve the quality in case of errors. Nevertheless in our implementation, for simplicity’s sake, we decided to discard the corrupted B and C bits and to take advantage of the motion vector information included in the class A bits to improve the concealment technique. In this case, the missing pixels are replaced with the ones in the previous frame as pointed to by the motion vectors. This method establishes a lower bound on the performance of an error resilient decoder, because the not checksummed bits are not used at all to improve the quality of the received video, even when parts of them are correct. An improved error resilient decoder could, for instance, decode the not checksummed bits until it detects an error, and then discard the remaining part of the packet.

We consider the ad hoc wireless scenario shown in Figure 3. A wireless node transmits a video sequence to a wireless receiver, by means of a variable number of relay nodes, in presence of other concurrent video traffic. The *ns* network simulator [14] is used to simulate the network behavior. The 802.11 network bandwidth is set to 11 Mb/s. Each video stream is sent using the IP/UDP/RTP protocol stack. The RTS/CTS mechanism has been enabled to limit the number of packet collisions due to the hidden node problem typical of ad hoc scenarios. Bit-level error traces collected from actual transmission experiments in different indoor environments are used to model the wireless channel behavior. The procedure to collect traces is the same used in [11]. The traces present a clear bursty behavior when channel conditions are poor.




Fig. 3. Wireless ad hoc transmission scenario. The actual number of relay nodes depends on the simulation. Each dashed ellipse represents the coverage of the device in its center, thus each end-to-end transmission needs to traverse all relay nodes.




Fig. 4. Packet loss and corruption rate for the various transmission schemes. *Carphone* sequence. Retry limit is equal to two.




Fig. 5. PSNR performance of the various transmission schemes, for the *carphone* sequence. Retry limit is equal to two.

Bursts are limited in length and include sporadically correct bits. Each time the simulator needs to send a packet, it loads from the trace an amount of bits equal to the packet size, and then, depending on the position of the errors in the packet, it determines the outcome of the packet transmission.

V. RESULTS

This section presents a simulative analysis of the proposed partial checksum technique. Results are presented for the *carphone* sequence, but similar results were obtained for the *foreman* sequence. In all the simulations, the retry limit is set to two.

Figure 4 shows the packet loss (PLR) and corruption (PCR) rate as a function of the number of relay nodes for different transmission schemes. The zero value implies that the wireless transmitter sends packets directly to the wireless receiver without using intermediate hops. Four additional concurrent video flows that use the same transmission scheme are present in the network. A packet is considered corrupted if it contains errors undetected by the checksum.

The lowest packet loss rate is offered by the header-only checksum scheme, but that scheme also presents the highest packet corruption rate. Packets are retransmitted, in fact, only

if errors affect the MAC or IP/UDP/RTP headers, a small fraction of the bits in the packets. The partial checksum technique presents a slightly higher packet loss rate but a much better packet corruption rate in comparison with the header-only checksum scheme. No erroneous packets are delivered by the full checksum scheme, at the expense of a higher packet loss rate, compared to the previous techniques. The packet loss rate of the unequal error protection technique is higher when compared with the other techniques because the value concerns both the important packets and the less important packets that are transmitted only once. The packet loss rate of class A packets (not shown) is very similar to the one of the partial checksum scheme. The strong increase when more than two hops are present is due to the severe network load caused by the transmission of two packets for each slice.

To assess the effectiveness of the proposed partial checksum scheme, the impact of the corrupted packets on video quality needs to be evaluated. Figure 5 compares the PSNR performance of the four different transmission schemes as a function of the number of relay nodes. Simulation conditions are the same as in Figure 4. The playout buffer is set to 500 ms. The performance of all the schemes clearly decreases as the number of relay nodes is increased, due to the higher packet




Fig. 6. Average number of transmissions for the various transmission schemes. *Carphone* sequence. The retry limit is equal to two.

loss rate. The proposed partial checksum scheme delivers a consistent gain, up to 1 dB, over the other transmission techniques, and in particular over the full checksum scheme. Hence the performance of the proposed partial checksum scheme confirms the expectation that the negative effect on quality due to corruption of the less sensitive bits in the packets is adequately counterbalanced by the lower number of losses. The unequal error protection scheme provides good performance when the transmission is carried out with no intermediate nodes, but the performance degrades significantly in comparison with the partial checksum scheme as the number of relay nodes increases. This is due to the overhead implied by transmitting twice the number of packets with respect to the other techniques. The ad hoc network, in fact, particularly suffers from the increased number of contentions caused by the higher number of packets. The header-only partial checksum scheme consistently delivers a lower performance with respect to all the other schemes, hence confirming that class A bits definitely need protection. Finally, note that the reported PSNR performance for the partial checksum scheme represents a lower bound on the performance achievable by the proposed scheme, since in our implementation the decoder does not exploit, for simplicity's sake, the correctly received parts of the B and C partitions contained in corrupted packets.

Figure 6 shows the average number of retransmissions per packet as a function of the number of relay nodes. Header-only checksum scheme excepted, the proposed partial checksum scheme presents the lowest value. A lower average number of retransmissions per packet causes fewer contentions for the channel, therefore the network throughput is improved, especially when the network is congested. The unequal error protection scheme presents a slightly higher average number of retransmissions than the proposed partial checksum technique. Note, however, that the graph only reports the value for class A packets. Considering the other type of packets as well, the total number of transmissions—hence contentions—nearly doubles in comparison with the other techniques, therefore the negative impact of the unequal error protection scheme on

the network load and throughput is much higher. The header-only checksum scheme provides very poor video quality performance, therefore it is pointless to argue about its impact on the network load and throughput.

VI. CONCLUSIONS

A partial checksum technique to enhance real-time video transmissions over 802.11-based ad hoc networks has been presented. We proposed to limit the coverage of the error detection mechanism to the most sensitive bit class contained in video packets, expecting that the negative effects of errors in the less sensitive bits should be more than adequately counterbalanced by the lower number of discarded video packets. The network load is reduced due to the lower number of retransmissions, with positive effects on all the transmissions. Simulation results obtained with *ns* in an ad hoc scenario with a varying number of relay nodes showed consistent video quality improvements (up to 1 dB PSNR) as well as network load reduction with respect to state-of-the-art techniques.

REFERENCES

- [1] "Wireless LAN medium access control (MAC) and physical layer (PHY) specifications," *ISO/IEC 8802-11, ANSI/IEEE Std 802.11*, 1999.
- [2] Q. Li and M. Van Der Schaar, "Providing adaptive QoS to layered video over wireless local area networks through real-time retry limit adaptation," *IEEE Transactions on Multimedia*, vol. 6, no. 2, pp. 278–290, April 2004.
- [3] S. Krishnamachari, M. Van Der Schaar, S. Choi, and X. Xu, "Video streaming over wireless LANs: A cross-layer approach," in *Proc. Packet Video Workshop*, Nantes, France, April 2003.
- [4] Y. Chen, J. C. Ye, C. R. Florjach, and K. Challopali, "Robust video streaming over wireless LAN with efficient scalable coding and prioritized adaptive transmission," in *Proc. IEEE Int. Conf. on Image Processing*, Barcelona, Spain, September 2003, vol. 3, pp. 285–288.
- [5] L.-A. Larzon, M. Degermark, S. Pink, L.-E. Jonsson, and G. Fairhurst, "The UDP-lite protocol," *draft-ietf-tsvwg-udp-lite-02.txt*, August 2003.
- [6] S. A. Khayam, S. Karande, H. Radha, and D. Loguinov, "Performance analysis and modeling of errors and losses over 802.11b LANs for high-bit-rate real-time multimedia," *Signal Processing: Image Communication*, vol. 18, no. 7, pp. 575–595, August 2003.
- [7] H. Zheng and J. Boyce, "An improved UDP protocol for video transmission over internet-to-wireless networks," *IEEE Trans. on Multimedia*, vol. 3, no. 3, pp. 356–365, September 2001.
- [8] G. Ding, H. Ghafoor, and B. Bhargava, "Error resilient video transmission over wireless networks," in *Proc. IEEE Workshop on Software Technologies for Future Embedded Systems*, Hakodate, Hokkaido, Japan, May 2003, pp. 31–34.
- [9] A. Servetti and J. C. De Martin, "Link-level unequal error detection for speech transmission over 802.11 wireless networks," in *Special Workshop in Maui (SWIM), Lectures by Masters in Speech Processing*, Maui, HI, January 2004.
- [10] H. Dong, I. D. Chakares, A. Gersho, E. Belding Royer, and J. D. Gibson, "Selective bit error checking at the MAC layer for voice over mobile ad hoc networks with IEEE 802.11," in *Proc. IEEE Wireless Communications and Networking Conference (WCNC)*, Atlanta, Georgia, March 2004.
- [11] E. Masala, M. Bottero, and J.C. De Martin, "Link level partial checksum for real time video transmission over 802.11 wireless networks," in *Proc. Packet Video Workshop*, Irvine, CA, December 2004.
- [12] Y. Wang and Q. Zhu, "Error control and concealment for video communication: a review," *Proceedings of the IEEE*, vol. 86, no. 5, pp. 974–997, May 1998.
- [13] ITU-T Rec. H.264 & ISO/IEC 14496-10 AVC, "Advanced video coding for generic audiovisual services," *ITU-T*, May 2003.
- [14] UCB/LBNL/VINT, "Network Simulator – ns – version 2," URL: <http://www.isi.edu/nsnam/ns>, 1997.